arXiv:1701.04102v1 [math.OC] 15 Jan 2017

Two-stage Linear Decision Rules for Multi-stage Stochastic Programming

Merve Bodur*[1] and James R. Luedtke†[2]

[1]Department of Mechanical and Industrial Engineering, University of Toronto
[2]Department of Industrial and Systems Engineering, University of Wisconsin-Madison

January 17, 2017

Abstract

Multi-stage stochastic linear programs (MSLPs) are notoriously hard to solve in general. Linear decision rules (LDRs) yield an approximation of an MSLP by restricting the decisions at each stage to be an affine function of the observed uncertain parameters. Finding an optimal LDR is a static optimization problem that provides an upper bound on the optimal value of the MSLP, and, under certain assumptions, can be formulated as an explicit linear program. Similarly, as proposed by Kuhn, Wiesemann, and Georghiou ("Primal and dual linear decision rules in stochastic and robust optimization" *Math. Program.* 130, 177–209, 2011) a lower bound for an MSLP can be obtained by restricting decisions in the dual of the MSLP to follow an LDR. We propose a new approximation approach for MSLPs, *two-stage LDRs*. The idea is to require only the state variables in an MSLP to follow an LDR, which is sufficient to obtain an approximation of an MSLP that is a *two-stage stochastic linear program* (2SLP). We similarly propose to apply LDR only to a subset of the variables in the dual of the MSLP, which yields a 2SLP approximation of the dual that provides a lower bound on the optimal value of the MSLP. Although solving the corresponding 2SLP approximations exactly is intractable in general, we investigate how approximate solution approaches that have been developed for solving 2SLP can be applied to solve these approximation problems, and derive statistical upper and lower bounds

*bodur@mie.utoronto.ca
†jim.luedtke@wisc.edu

on the optimal value of the MSLP. In addition to potentially yielding better policies and bounds, this approach requires many fewer assumptions than are required to obtain an explicit reformulation when using the standard static LDR approach. As an illustrative example we apply our approach to a capacity expansion model, and find that the two-stage LDR policy has expected cost between 20% and 34% lower than the static LDR policy, and in the dual yields lower bounds that are between 0.1% and 3.3% better.

Keywords: Multi-stage stochastic programming, linear decision rules, two-stage approximation

1 Introduction

We present a new approach for approximately solving multi-stage stochastic linear programs (MSLPs). MSLPs model dynamic decision-making processes in which a decision is made, a stochastic outcome is observed, another decision is made, and so on, for T stages. At each stage, the decision vectors are constrained by linear constraints that depend on the history of observed stochastic outcomes. A solution of MSLP is a policy, which defines the decisions to be made at each stage as a function of the observed outcomes up to that stage. The objective in an MSLP is to choose a policy that minimizes the expected cost over all stages. Although MSLPs can be used to model a wide variety of problems (e.g., [52]), they are notoriously hard to solve in general [15, 50].

There are a variety of methods available for MSLPs in the case that the stochastic process is represented by a scenario tree [11, 22, 26, 48]. Such algorithms include nested Benders decomposition [8, 9, 19], progressive hedging [42], and aggregation and partitioning [2, 7], and enable the solution of MSLPs with possibly very large scenario trees. Unfortunately, as discussed in [50], the size of a scenario tree needed to obtain even a modestly accurate approximation grows exponentially in the number of stages. For example, a 10-stage problem in which the uncertainty in each stage is represented by just 50 realizations would yield a scenario tree having nearly $2 \cdot 10^{15}$ scenarios, making any approach that requires even a single pass through the scenario tree impossible.

Under some conditions, including stage-wise independence of the random variables, stochastic dual dynamic programming (SDDP) [36] can overcome the difficulty in exploding scenario tree size, by constructing a single value function approximation for each stage.

The SDDP algorithm converges almost surely on a finite scenario tree [13, 46] (see also [32, 38] for related results). In some cases, such as additive dependence [27] and Markov dependence [37], the assumption of stage-wise independence can be satisfied via an appropriate reformulation, e.g., see Example 10 in [47]. However, such reformulations are not applicable for stage-wise dependence of random recourse matrices (i.e., in the coefficients of the constraints). Similar approaches that exploit stage-wise independence and value function approximations include multi-stage stochastic decomposition [44] and approximate dynamic programming [40].

An alternative approach to handling the complexity of MSLP is to restrict the functional form of the policy. One such approach is the use of *linear decision rules* (LDRs). The idea of an LDR is to require that all decisions made in each stage be a linear (or affine) function of the observed random outcomes up to that stage. The problem then reduces to a static problem of finding the best LDR, whose expected cost then yields an upper bound on the optimal value of the MSLP. In this paper, we refer this use of LDR as a *static LDR*. While LDRs have a long history (see, e.g., [18]), they have recently gained renewed interest in the mathematical optimization literature after their application to adjustable robust optimization in [4]. The adaptation of this approach to MSLP was presented in [50], and Kuhn et al. [29] analyzed the application of a static LDR in the dual of the MSLP, which then yields a lower bound on the optimal value of the MSLP. Moreover, under certain assumptions, the static approximations obtained after restricting the primal and dual policies to be LDR are both tractable linear programs, as shown in [50] and [29], respectively. The assumptions include stage-wise independence (or a slight generalization), compact and polyhedral support, and fixed recourse (i.e., uncertainty is limited to the right-hand side of the constraints). While in some cases static LDR policies provide high quality approximations to MSLP, they have potential to be significantly suboptimal. Better policies (primal and dual) can be obtained by considering more flexible (nonlinear) rules such as (static) piecewise linear decision rules [12] and polynomial decision rules [3].

We propose a new use of LDR, which we refer to as *two-stage linear decision rules*. The key idea is to partition the decision variables into *state* and *recourse* decision variables, with the property that if the state variables are fixed, then the problem decouples into a separate problem for each stage, involving only recourse decision variables. If one applies an LDR *only* to the state variables, then the problem reduces to a *two-stage* stochastic linear program (2SLP), in contrast to a *static* problem which is obtained when using a static LDR. The advantage of two-stage LDRs is that they free the recourse variables from

the LDR requirement, thus allowing for a potentially improved policy. Indeed, there exist feasible 2SLPs that are *infeasible* if one enforces LDR on the recourse variables [18]. This idea of reducing an MSLP to a 2SLP is similar to that proposed by Ahmed [1], except that in [1] the state variables are completely decided in the first-stage and fixed, whereas we allow them to vary according to an LDR. We also consider applying two-stage LDR in the dual of an MSLP, exploiting the observation that imposing an LDR restriction *only* on the dual variables associated with the *state equations* is sufficient to obtain a 2SLP that approximates the multi-stage dual problem. We investigate how approximate solutions to the associated primal and dual approximation problems can be used to obtain feasible policies with associated statistical estimates on the optimality gap. Our analysis suggests that this can be done under mild assumptions, for example that the primal problem exhibits relatively complete recourse (i.e., for any current state there exists a feasible next decision and state) and has a bounded feasible region with probability 1. We illustrate the two-stage LDR approach on a capacity expansion problem proposed in [14], and find that, for this problem, two-stage LDR yields significantly better primal policies (upper bounds), and modestly improves on the lower bounds, when compared to static LDR.

A significant challenge to using two-stage LDR is that the resulting 2SLP is in general intractable to solve exactly. Indeed, 2SLP is $\#P$-hard [15, 21] due to the difficulty in evaluating the expectation of the recourse function. However, as argued in [50], under mild conditions Monte Carlo sampling-based methods can provide solutions of modest accuracy to a 2SLP (such a statement cannot be made for MSLP). Thus, an important benefit of the two-stage LDR approach is that it enables the application of the long history of research into solving 2SLPs to the multi-stage setting. While using a sampling-based method may lead to a suboptimal solution of the 2SLP approximations, our hope is that this suboptimality may be more than offset by the improvement gained by eliminating the LDR requirement on the recourse decisions that is imposed when using a static LDR. In addition, when using a sampling-based approach, the assumptions that are required for obtaining a tractable reformulation when applying static LDR are no longer needed. In particular, the random variables need not have polyhedral (or even bounded) support, the constraint matrices may be random and dependent across time stages, and the LDR may be based on nonlinear functions of the random variables.

The two-stage LDR approach we propose can also be applied to certain multi-stage stochastic *integer programs*, in which some of the decision variables are required to be integer valued. In particular, for the primal problem, the approach applies directly provided

integrality restrictions are imposed only on the recourse variables. When the state variables have integrality restrictions as well, the form of the decision rule applied to the state variables must be modified, but the two-stage approach still applies. We refer to [6] for one possible such decision rule based on piecewise-linear binary functions. We remark that combining our approach with that of [6] would have potential benefit in terms of both tractability and policy quality, as removing the piecewise-linear binary decision rule requirement from the recourse variables both eliminates the need to design such a rule, and gives those decisions more flexibility.

The rest of this paper is organized as follows. Section 2 defines the MSLP, reviews the static LDR approach, and presents the proposed two-stage LDR approach, including discussion of how to solve the approximate problem and obtain statistical upper bounds on the original MSLP. Section 3 conducts a similar analysis for the dual of an MSLP, yielding an approach for finding statistical lower bounds on an MSLP. We present an illustrative application in Section 4 and make concluding remarks in Section 5.

2 Primal two-stage linear decision rules

We formulate an MSLP with $T \geq 2$ stages as follows, where throughout the paper, for integers $a \leq b$, $[a, b] := \{a, a+1, \ldots, b\}$ and $[b] := \{1, \ldots, b\}$:

$$\min_{x,s} \ \mathbb{E}\Big[\sum_{t\in[T]} c_t(\xi^t)^\top x_t(\xi^t) + h_t(\xi^t)^\top s_t(\xi^t) \Big] \tag{1a}$$

$$\text{s.t.} \ \ A_t(\xi^t)s_t(\xi^t) + B_t(\xi^t)s_{t-1}(\xi^{t-1}) + C_t(\xi^t)x_t(\xi^t) = b_t(\xi^t), \quad t \in [T], \ \mathbb{P}\text{-a.s.}, \tag{1b}$$

$$D_t(\xi^t)s_t(\xi^t) + E_t(\xi^t)x_t(\xi^t) \geq d_t(\xi^t), \qquad t \in [T], \ \mathbb{P}\text{-a.s.}, \tag{1c}$$

$$x_t(\xi^t) \in \mathbb{R}^{p_t}, \ s_t(\xi^t) \in \mathbb{R}^{q_t}, \qquad t \in [T], \ \mathbb{P}\text{-a.s.} \tag{1d}$$

Here, $\{\xi_t\}_{t=1}^T$ is a stochastic process with probability distribution $\mathbb{P}$ and support Ξ, where $\xi_1 = 1$ for all $\xi \in \Xi$ (i.e., data in stage 1 is deterministic), ξ_t is a random vector taking values in $\mathbb{R}^{\ell_r}$ for $t \in [2, T]$, and $\xi^t := (\xi_1, \ldots, \xi_t)$ for $t \in [T]$. We denote $\ell^t := \sum_{r=1}^t \ell_r$ for $t \in [T]$. The s and x variables are referred to as *state* and *recourse* variables, respectively. Similarly, (1b) and (1c) are referred to as *state equations* and *recourse constraints*, respectively. The objective is to minimize the expected total cost. The constraints are required to be almost surely satisfied with respect to the distribution of the stochastic process, denoted by "$\mathbb{P}$-a.s." We note that any MSLP can be brought into the form of (1) by introducing additional

variables and constraints. Throughout the paper, we assume that (1) is feasible and has an optimal solution, and we denote its optimal objective value as z^{MSLP}.

2.1 Static linear decision rules

A tractable approximation of MSLP can be obtained by restricting the decision policies to a certain form, i.e., by restricting the decisions to be a special function of the uncertain parameters. A linear decision rule is a policy in which the decisions at each stage t are restricted to be a linear function of the observed random variables ξ^t up that stage. We refer to the policies in which all the decisions are required to follow an LDR as *static LDR policies*. Specifically, a static LDR policy has the form:

$$s_t(\xi^t) = \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\beta_{tk}, \tag{2a}$$

$$x_t(\xi^t) = \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\theta_{tk}, \tag{2b}$$

where $\theta_{tk} \in \mathbb{R}^{p_t}$ and $\beta_{tk} \in \mathbb{R}^{q_t}$ are free parameters of the LDR, and $\Phi_{tk}(\xi^t) : \mathbb{R}^{\ell^t} \to \mathbb{R}$ for all $t \in [T]$ and $k \in [K_t]$ are given *LDR basis functions*. We refer to the θ and β variables as the *LDR variables*. We assume $K_1 = 1$ and $\Phi_{t1}(\xi^t) = 1$ for $t \in [T]$. Often, the basis functions are the uncertain parameters themselves, i.e., $K_t = \ell^t$ and $\Phi_{tk}(\xi^t) = (\xi^t)_k$, where $(\xi^t)_k$ denotes the k^{th} component of ξ^t vector. In this case, we refer to the basis functions as the *standard basis functions*. Note that the convention $\xi_1 \equiv 1$ implies that the decisions made at stage t are actually *affine* functions of the random variables $(\xi_2, \ldots, \xi_t)$.

Substituting the LDRs of the form (2) into MSLP given in (1) yields the following approximation of MSLP, which we call LDR-LP:

$$\begin{aligned}
\min_{\theta,\beta}\ & \mathbb{E}\Big[\sum_{t\in[T]}\sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\big(c_t(\xi^t)^\top\theta_{tk} + h_t(\xi^t)^\top\beta_{tk}\big)\Big] \\
\text{s.t.}\ & \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\big(A_t(\xi^t)\beta_{tk} + C_t(\xi^t)\theta_{tk}\big) + \sum_{k\in[K_{t-1}]} \Phi_{t-1,k}(\xi^{t-1})B_t(\xi^t)\beta_{t-1,k} = b_t(\xi^t),\ t\in[T],\ \mathbb{P}\text{-a.s.}, \\
& \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\big(D_t(\xi^t)\beta_{tk} + E_t(\xi^t)\theta_{tk}\big) \geq d_t(\xi^t), \qquad t\in[T],\ \mathbb{P}\text{-a.s.}, \\
& \theta_{tk} \in \mathbb{R}^{p_t},\ \beta_{tk} \in \mathbb{R}^{q_t}, \qquad t\in[T],\ k\in[K_t].
\end{aligned}$$

We let z^{LDR} denote the optimal value of LDR-LP, where here and elsewhere, we adopt the convention that if a minimization (maximization) problem is infeasible, the associated optimal value is defined to be $+\infty$ ($-\infty$). Note that all the decision variables are deterministic, i.e., they have to be determined before observing any random outcomes, and hence this problem is a static problem. LDR-LP is a semi-infinite program having infinitely many constraints. It is observed in [50] (see also [12, 29]) that LDR-LP can be reformulated as a linear program (LP) using robust optimization techniques under the following assumptions:

A1. The standard basis functions are used.

A2. Fixed recourse: the constraint matrices, A_t, B_t, C_t, D_t, E_t, for all $t \in [T]$, are independent of the random vector ξ.

A3. The conditional expectation $\mathbb{E}(\xi|\xi^t)$ is almost surely linear in ξ^t for all $t \in [T]$ (e.g., this occurs when $\{\xi_t\}_{t\in[T]}$ are mutually independent, known as *stage-wise independence*).

A4. The support, Ξ, is a nonempty compact polyhedron.

The LP reformulation has constraints of the form $(\beta, \theta, w) \in \mathcal{D}$, where w are auxiliary variables and $\mathcal{D}$ is an explicitly given polyhedron. The size of this LP scales well (typically grows only quadratically) with the number of stages T. Moreover, the LP does not require any discretization of $\mathbb{P}$ (e.g., by Monte Carlo sampling), and instead only uses a polyhedral description of Ξ and the second order moment matrix of the random variables. These results have been generalized in [20] to the case of conic support, where A4 is replaced by an assumption that Ξ is described by a finite set of conic inequalities, in which case LDR-LP is reformulated as a conic program.

As LDR-LP is a restriction of MSLP, it provides an upper bound to MSLP. However, the benefit of tractability comes at the expense of loss of optimality. That is, the obtained upper bound can be substantially far from the optimal value of MSLP. Indeed, for 2SLPs, the optimal recourse decisions are very rarely linear in the random variables, but there always exists an optimal *piecewise linear* decision rule [18].

2.2 Two-stage linear decision rules

We propose *two-stage LDRs* which yield upper bounds to MSLP that cannot be worse than the ones obtained by LDR-LP. The key idea is to apply LDR only on the state

variables to obtain a *two-stage* approximation of MSLP, rather than a static approximation. Substituting the LDRs of the form (2a) into the MSLP given in (1) yields

$$
\begin{aligned}
\min_{x,\beta}\ & \sum_{t\in[T]} \mathbb{E}\big[c_t(\xi^t)^\top x_t(\xi^t)\big] + \sum_{t\in[T]}\sum_{k\in[K_t]} \mathbb{E}\big[h_t(\xi^t)^\top \Phi_{tk}(\xi^t)\beta_{tk}\big] \\
\text{s.t. } & A_t(\xi^t)\sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\beta_{tk} + C_t(\xi^t)x_t(\xi^t) + B_t(\xi^t)\sum_{k\in[K_{t-1}]} \Phi_{t-1,k}(\xi^{t-1})\beta_{t-1,k} = b_t(\xi^t),\ t\in[T],\ \mathbb{P}\text{-a.s.}, \\
& D_t(\xi^t)\sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\beta_{tk} + E_t(\xi^t)x_t(\xi^t) \geq d_t(\xi^t), && t\in[T],\ \mathbb{P}\text{-a.s.}, \\
& x_t(\xi^t)\in\mathbb{R}^{p_t}, && t\in[T],\ \mathbb{P}\text{-a.s.}, \\
& \beta_{tk}\in\mathbb{R}^{q_t}, && t\in[T],\ k\in[K_t].
\end{aligned}
$$

We denote this problem as LDR-2SLP and the optimal value of this problem as z^{2S}. This problem is a 2SLP, which can be equivalently written as follows, where we drop dependence of the first-stage variables on $\xi^1 \equiv 1$ and use $\beta_{11} = s_1(\xi^1) = \Phi_{11}(\xi^1)\beta_{11}$,

$$
\begin{aligned}
z^{2S} = \min_{x_1,\beta}\ & c_1^\top x_1 + \sum_{t\in[T]}\sum_{k\in[K_t]} \mathbb{E}\big[\Phi_{tk}(\xi^t)h_t(\xi^t)^\top\big]\beta_{tk} + \mathbb{E}[F(\beta,\xi)] \\
\text{s.t. } & A_1\beta_{11} + C_1x_1 = b_1, \\
& D_1\beta_{11} + E_1x_1 \geq d_1, \\
& x_1\in\mathbb{R}^{p_1}, \\
& \beta_{tk}\in\mathbb{R}^{q_t},\ t\in[T], k\in[K_t],
\end{aligned}
$$

where $F(\beta,\xi) := \sum_{t\in[2,T]} f_t(\beta,\xi^t)$ and for $t\in[2,T]$, $f_t(\beta,\xi^t)$ is defined to be the optimal objective value of

$$
\begin{aligned}
\min_{x_t}\ & c_t(\xi^t)^\top x_t && \text{(4a)} \\
\text{s.t. } & C_t(\xi^t)x_t = b_t(\xi^t) - A_t(\xi^t)\sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\beta_{tk} - B_t(\xi^t)\sum_{k\in[K_{t-1}]} \Phi_{t-1,k}(\xi^{t-1})\beta_{t-1,k}, && \text{(4b)} \\
& E_t(\xi^t)x_t \geq d_t(\xi^t) - D_t(\xi^t)\sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\beta_{tk}, && \text{(4c)} \\
& x_t\in\mathbb{R}^{p_t}. && \text{(4d)}
\end{aligned}
$$

The following proposition, immediate from the definitions of the associated problems, summarizes the relationship between the optimal values of MSLP, LDR-LP, and LDR-2SLP.

Proposition 2.1. *The following inequalities hold:*

$$z^{MSLP} \leq z^{2S} \leq z^{LDR}.$$

The difference between z^{LDR} and z^{2S} can be arbitrarily large. In particular, an example is given in [18] of a 2SLP having relatively complete recourse for which LDR-LP is infeasible ($z^{LDR} = \infty$), while the 2SLP (and hence LDR-2SLP) is feasible.

Unfortunately, the techniques used to derive a static approximation of LDR-LP do not yield an efficiently computable reformulation of LDR-2SLP, even under assumptions A1-A4. In the next section, we review approaches for obtaining an approximate solution, say $\hat{\beta}$, of LDR-2SLP. Then, in Section 2.4 we discuss techniques for estimating an upper bound on z^{MSLP} using such a solution.

2.3 Approximate solution of LDR-2SLP

There is a huge literature on (approximately) solving 2SLP problems. In this section, we present a brief overview of relevant approaches, with a focus on identifying the required assumptions. We refer the reader to [48] for more details.

A common approach for approximately solving a 2SLP is *sample average approximation*, in which $\mathbb{P}$ is approximated by a discrete probability measure $\hat{\mathbb{P}}$ that assigns positive weights only to a finite (relatively small) number of realizations of ξ which are called *scenarios*. In this way, the intractable expectation term is replaced with a sum. Scenarios may be constructed by a variety of techniques, such as Monte Carlo, quasi-Monte Carlo, and Latin hypercube sampling (e.g., [24, 28, 31, 35, 49]). For the purpose of this paper, we consider only the conceptually simplest case in which scenarios are generated via independent Monte Carlo sampling.

Let $\xi_j, j = 1, \ldots, N$, be an independent and identically distributed (i.i.d.) random sample of the random vector ξ, and define the sample average approximation (SAA) problem:

$$\hat{z}_N^{2S} := \min_{x_1,\beta} \; c_1^\top x_1 + \sum_{t\in[T]} \sum_{k\in[K_t]} \mathbb{E}\big[\Phi_{tk}(\xi^t) h_t(\xi^t)^\top\big]\beta_{tk} + \frac{1}{N}\sum_{j\in[N]} F(\beta,\xi_j) \tag{5a}$$

$$\text{s.t.} \;\; A_1\beta_{1,1} + C_1 x_1 = b_1, \tag{5b}$$

$$D_1\beta_{1,1} + E_1 x_1 \geq d_1, \tag{5c}$$

$$x_1 \in \mathbb{R}^{p_1}, \tag{5d}$$

$$\beta_{tk} \in \mathbb{R}^{q_t}, \; t \in [T], k \in [K_t]. \tag{5e}$$

Once the sample is fixed, the SAA problem can be solved by any approach for solving the above, now deterministic, problem. In particular the L-shaped decomposition algorithm [51] or a regularized variant [30, 43] can be applied, with the further advantage that the subproblem obtained with fixed $\hat{\beta}$ decomposes by both scenario and stage due to the relationship $F(\beta,\xi) = \sum_{t\in[2,T]} f_t(\beta,\xi^t)$.

If there exists a $\bar{\beta}$ such that $\mathbb{E}[F(\beta,\xi)] < \infty$ for all β in a neighborhood of $\bar{\beta}$, then because $F(\cdot,\xi)$ is a convex function for all $\xi \in \Xi$, Theorem 5.4 of [48] applies and implies that

$$\hat{z}_N^{2S} \to z^{2S} \quad \text{with probability 1 as } N \to \infty \tag{6}$$

and also that the set of optimal solutions to (5) converges to the set of optimal solutions of LDR-2SLP.

The problem (5) can also be used to obtain a statistical confidence interval on the optimal value z^{2S} of LDR-2SLP. A basic result in this development is that $\mathbb{E}[\hat{z}_N^{2S}] \leq z^{2S}$; that is, on average, the optimal value of the sample average approximation provides a lower bound on z^{2S}. A confidence interval on $\mathbb{E}[\hat{z}_N^{2S}]$ can be obtained by solving (5) multiple times, with different independent samples, and thus a lower bound on z^{2S} can be estimated [33]. Unfortunately, this does not translate into a lower bound on the true quantity of interest, z^{MSLP}, since $z^{2S} \geq z^{MSLP}$. We return to the question of obtaining a lower bound on z^{MSLP} in Section 3.

Stronger results, in particular addressing the rate of convergence of $\hat{z}_N^{2S}$ to z^{2S}, require additional assumptions. An important assumption is *relatively complete recourse*, which states that for any feasible first-stage decision β, the expected value of the second-stage value function $\mathbb{E}[F(\beta,\xi)]$ is finite, which occurs if and only if $F(\beta,\xi) < +\infty$ $\mathbb{P}$-a.s. As LDR-2SLP may not necessarily satisfy this property, we consider the possibility of adding an

additional set of constraints, $\beta \in \mathcal{B} \subseteq \mathbb{R}^{\tau_P}$ where $\tau_P := \sum_{t\in[T]} q_t K_t$, to the problem LDR-2SLP (and to the SAA (5)). We now state the relatively complete recourse assumption for $\mathcal{B}$.

Assumption 2.2. *For all $\beta \in \mathcal{B}$, $F(\beta, \xi) < +\infty$ $\mathbb{P}$-a.s..*

If LDR-2SLP already has relatively complete recourse, then we can take $\mathcal{B} = \mathbb{R}^{\tau_P}$, and hence impose no additional constraints. Otherwise, adding the constraints $\beta \in \mathcal{B}$ has the potential to make the approximation more conservative. Derivation of a set $\mathcal{B}$ that satisfies Assumption 2.2 is a difficult task in general. However, relatively complete recourse can often be achieved by appropriate modeling, e.g., by introducing "artificial" variables that allow violation of a constraint, where the violation amount is then penalized in the objective function. Derivation of a set $\mathcal{B}$ satisfying Assumption 2.2 may then be possible using ad hoc techniques. We provide one such example in our illustrative example in Section 4 in which $\mathcal{B} = [0, M]^{\tau_P}$ for suitably chosen M is easily checked to satisfy Assumption 2.2. Another possibility, if assumptions A1-A4 hold, is to use the robust optimization techniques used in [29, 50] to derive a tractable set $\mathcal{D}$ such that (β, θ) is feasible to LDR-LP if and only if there are values of auxiliary variables w such that $(\beta, \theta, w) \in \mathcal{D}$. Then, $\mathcal{B} = \text{proj}_\beta(\mathcal{D})$ would satisfy Assumption 2.2. This construction of $\mathcal{B}$ is more conservative than necessary, because it restricts β to values for which there is also a θ that makes the static LDR policy defined in (2) feasible $\mathbb{P}$-a.s.. However, the resulting policy could still potentially be better (and for sure would not be worse) than the static LDR policy obtained from LDR-LP, since enforcing $\beta \in \text{proj}_\beta(\mathcal{D})$ would not require the recourse decisions to follow an LDR policy (it only requires existence of a feasible LDR policy).

Under Assumption 2.2 and additional mild assumptions, estimates for the sample size N required to ensure that an optimal solution of (5) is an ϵ-optimal solution to LDR-2SLP with probability at least $1 - \alpha$ are obtained in [45, 49, 50]. These bounds scale linearly with the dimension of the first-stage variables, β and x_1 in this case. Dependence on the confidence α is $\ln(1/\alpha)$ so that high confidence can be achieved, but the dependence on ϵ is $O(1/\epsilon^2)$, which is why sampling is limited to obtaining "medium accuracy" solutions.

LDR-2SLP, with the additional constraints $\beta \in \mathcal{B}$, can also be approximately solved by stochastic approximation [41] or one of its robust extensions, e.g., [34, 39], when Assumption 2.2 holds.

Finally, we remark that if we cannot derive a set $\mathcal{B}$ satisfying Assumption 2.2, results about sampling-based approximation of chance-constrained programs derived in [10] can

be used to show that an optimal solution of the SAA problem (5) yields a policy that is feasible for a large fraction of the random outcomes. This has been previously used in [5, 6] when using sampling to approximately solve static approximations derived from finitely adaptable and piecewise-linear decision rules. Although the two-stage LDR policy itself is not necessarily feasible $\mathbb{P}$-a.s. in this case, in the next section we discuss how an approximate solution $\hat{\beta}$ could still be used to guide a feasible policy.

2.4 Feasible policies and upper bounds on z^{MSLP}

Let $(\hat{x}_1, \hat{\beta})$ be an approximate first-stage solution to LDR-2SLP. We discuss how such a solution can be used to estimate upper bounds on z^{MSLP}.

We first assume that $\hat{\beta} \in \mathcal{B}$ for a set $\mathcal{B}$ satisfying Assumption 2.2. In this case, $(\hat{x}_1, \hat{\beta})$ defines a feasible solution to LDR-2SLP, and a feasible two-stage LDR policy for MSLP. Thus, the expected cost of this solution provides an upper bound on z^{2S} and z^{MSLP}. The expected cost of the policy defined by $(\hat{x}_1, \hat{\beta})$ can be estimated by generating an independent sample of ξ, say $\{\xi_j\}_{j=1}^{N'}$, and computing

$$c_1^\top \hat{x}_1 + \sum_{t\in[T]} \sum_{k\in[K_t]} \mathbb{E}\big[\Phi_{tk}(\xi^t) h_t(\xi^t)^\top\big] \hat{\beta}_{tk} + \frac{1}{N'} \sum_{j\in[N']} F(\hat{\beta}, \xi_j).$$

(The second term in the objective can also be estimated by sampling if $\mathbb{E}[\Phi_{tk}(\xi^t)h_t(\xi^t)^\top]$ cannot be calculated efficiently.) Because $\hat{\beta}$ is fixed in this evaluation step, it would generally be computationally feasible to use $N' \gg N$. The values $F(\hat{\beta}, \xi_j)$ for $j \in [N']$ can also be used to construct a confidence interval on the objective value of $(\hat{x}_1, \hat{\beta})$, and hence a statistical upper bound on z^{MSLP}.

We next consider the case when we do not know $\hat{\beta} \in \mathcal{B}$ for a set $\mathcal{B}$ satisfying Assumption 2.2, so that we do not know a priori that $\mathbb{E}[F(\hat{\beta}, \xi)]$ is finite. To construct a policy in this case, we make the following relatively complete recourse assumption for the original problem MSLP.

Assumption 2.3. *For all $\xi \in \Xi$, and each $t \in [2, T]$, if the random vectors $\{(s_{t'}(\xi^{t'}), x_{t'}(\xi^{t'})\}_{t'\in[t-1]}$ satisfy the constraints of MSLP for $t' \in [t-1]$, then there exists (s_t, x_t) that satisfies the*

constraints of MSLP in stage t:

$$A_t(\xi^t)s_t + C_t(\xi^t)x_t = b_t(\xi^t) - B_t(\xi^t)s_{t-1}(\xi^{t-1})$$
$$D_t(\xi^t)s_t + E_t(\xi^t)x_t \geq d_t(\xi^t).$$

In other words, this assumption states that in any stage t, for any value of the previous state variables $s_{t-1}(\xi^{t-1})$ that could be obtained from past realizations of the random outcomes and past feasible decisions, there always exists a feasible set of decisions in the current stage (see e.g., [23]).

When we do not have $\hat{\beta} \in \mathcal{B}$ for some $\mathcal{B}$ satisfying Asumption 2.2, it may not be feasible to implement two-stage LDR based on $\hat{\beta}$ directly. However, under Assumption 2.3, we can implement a policy which is guided by $\hat{\beta}$, which we refer to as a *state-target tracking (STT) policy*. Specifically, at stage $t = 1$, we implement the solution $x_1^{STT} = \hat{x}_1$ and $s_1^{STT} = \hat{\beta}_{11}$. Then, for each stage $t \in [2, T]$, we first observe ξ_t (thus, we have ξ^t), and define $\hat{s}_t(\hat{\beta}, \xi^t) := \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\hat{\beta}_{tk}$ as the ending state that would be obtained if the LDR policy $\hat{\beta}$ were used to determine $s_t(\xi^t)$. We then solve the problem (deterministic for this fixed ξ^t):

$$\min_{x_t, s_t} \quad c_t(\xi^t)^\top x_t + h_t(\xi^t)^\top s_t + \rho\|s_t - \hat{s}_t(\hat{\beta}, \xi^t)\| \tag{7a}$$
$$\text{s.t.} \quad A_t(\xi^t)s_t + C_t(\xi^t)x_t = b_t(\xi^t) - B_t(\xi^t)s_{t-1}^{STT}(\xi^{t-1}),$$
$$D_t(\xi^t)s_t + E_t(\xi^t)x_t \geq d_t(\xi^t), \tag{7b}$$
$$x_t(\xi^t) \in \mathbb{R}^{p_t}, s_t(\xi^t) \in \mathbb{R}^{q_t}, \tag{7c}$$

where $\rho \geq 0$ is a parameter of the policy and $\|\cdot\|$ is any norm, and let the solution be $x_t^{STT}(\xi^t), s_t^{STT}(\xi^t)$. For any $\xi \in \Xi$, all problems in this sequence are feasible when Assumption 2.3 holds, and hence this yields a feasible policy to MSLP. Observe that when $\rho = 0$, the policy reduces to a pure myopic policy that only considers the cost of decisions in each stage, without considering the impact of s_t on future costs. Using larger values of $\rho > 0$ has the effect of encouraging the decisions to be made in a way that keeps the state close to what would have been achieved if we could exactly follow the LDR policy defined by $\hat{\beta}$ on the state variables. The cost of the STT policy under a sample ξ of the stochastic process is

$$\sum_{t\in[T]} \big(c_t(\xi^t)^\top x_t^{STT}(\xi^t) + h_t(\xi^t)^\top s_t^{STT}(\xi^t)\big).$$

The expected cost of the STT policy is an upper bound on the optimal value of MSLP, and a confidence interval on this expected cost can be obtained by simulation with independent replications. The value of the parameter ρ can be heuristically selected by estimating the expected cost of the policy under different values of ρ and choosing the most promising value, or by using optimization via simulation techniques [17, 25]. Note that using the STT policy, even the decisions $s_t^{STT}(\xi^t)$ may not necessarily have the form of an LDR. Thus, simulating this policy yields an estimate of an upper bound on z^{MSLP}, but not necessarily on z^{2S}.

3 Dual two-stage linear decision rules

In this section, we apply two-stage LDR to the dual of MSLP, with the goal of obtaining lower bounds on the optimal value of MSLP. The dual of MSLP, which we refer to as D-MSLP, is the problem (see [16]):

$$\max_{\lambda,\gamma} \ \mathbb{E}\Big[\sum_{t\in[T]} b_t(\xi^t)^\top \lambda_t(\xi^t) + d_t(\xi^t)^\top \gamma_t(\xi^t)\Big] \tag{8a}$$

$$\text{s.t.} \ \mathbb{E}\Big[B_{t+1}(\xi^{t+1})^\top \lambda_{t+1}(\xi^{t+1})\Big|\xi^t\Big]$$

$$+ A_t(\xi^t)^\top \lambda_t(\xi^t) + D_t(\xi^t)^\top \gamma_t(\xi^t) = h_t(\xi^t), \quad t\in[T],\ \mathbb{P}\text{-a.s.}, \tag{8b}$$

$$C_t(\xi^t)^\top \lambda_t(\xi^t) + E_t(\xi^t)^\top \gamma_t(\xi^t) = c_t(\xi^t), \quad t\in[T],\ \mathbb{P}\text{-a.s.}, \tag{8c}$$

$$\gamma_t(\xi^t) \geq 0, \quad t\in[T],\ \mathbb{P}\text{-a.s.}, \tag{8d}$$

$$\lambda_t(\xi^t)\in\mathbb{R}^{m_t},\ \gamma_t(\xi^t)\in\mathbb{R}^{n_t}, \quad t\in[T],\ \mathbb{P}\text{-a.s.}, \tag{8e}$$

where $B_{T+1}(\xi^{T+1}) = 0$. For $t\in[T]$, the dual decisions $\lambda_t(\cdot)$ (corresponding to constraints (1b) in MSLP) and $\gamma_t(\cdot)$ (corresponding to constraints (1c) in MSLP) are functions of the data ξ^t observed up to stage t. Weak duality holds for MSLP and D-MSLP, i.e., the optimal objective value of D-MSLP provides a lower bound on z^{MSLP}. Moreover, under some conditions, strong duality holds (i.e., optimal value of D-MSLP equals z^{MSLP}) [16], although we only require weak duality.

3.1 Static linear decision rules

In [29] it has been proposed to use static LDR to obtain a tractable approximation of D-MSLP, and thus an efficiently computable lower bound on z^{MSLP}. Specifically, the idea

is to require all the dual decisions to be LDR, i.e.,

$$\lambda_t(\xi^t) = \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\alpha_{tk}, \tag{9a}$$

$$\gamma_t(\xi^t) = \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\pi_{tk}, \tag{9b}$$

where $\alpha_{tk} \in \mathbb{R}^{m_t}, \pi_{tk} \in \mathbb{R}^{n_t}$, for all $t \in [T]$ and $k \in [K_t]$ are the parameters of the decision rule. Imposing (9) yields the following static approximation of D-MSLP, which we call D-LDR-LP:

$$\begin{aligned}
\max_{\alpha,\pi}\ & \mathbb{E}\Big[\sum_{t\in[T]}\sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\big(b_t(\xi^t)^\top\alpha_{tk} + d_t(\xi^t)^\top\pi_{tk}\big)\Big] \\
\text{s.t. }\ & \mathbb{E}\Big[B_{t+1}(\xi^{t+1})^\top \sum_{k\in[K_{t+1}]} \Phi_{tk}(\xi^{t+1})\alpha_{t+1,k}\Big|\xi^t\Big] \\
& \quad + \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\big(A_t(\xi^t)^\top\alpha_{tk} + D_t(\xi^t)^\top\pi_{tk}\big) = h_t(\xi^t),\ t\in[T],\ \mathbb{P}\text{-a.s.}, \\
& \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\big(C_t(\xi^t)^\top\alpha_{tk} + E_t(\xi^t)^\top\pi_{tk}\big) = c_t(\xi^t), \qquad t\in[T],\ \mathbb{P}\text{-a.s.}, \\
& \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\pi_{tk} \geq 0, \qquad t\in[T],\ \mathbb{P}\text{-a.s.}, \\
& \alpha_{tk}\in\mathbb{R}^{m_t},\ \pi_{tk}\in\mathbb{R}^{n_t}, \qquad t\in[T],\ k\in[K_t].
\end{aligned}$$

We refer to the optimal value of D-LDR-LP as v^{LDR}. The semi-infinite program D-LDR-LP can be reformulated as an efficiently solvable LP if assumptions A1 - A4 stated in Section 2 hold, and also the problem MSLP is strictly feasible [29]. If assumption A4 is replaced by an assumption that Ξ is described by conic inequalities, D-LDR-LP can be reformulated as a conic program [20].

3.2 Two-stage linear decision rules

Examining the structure of D-MSLP, given in (8), we observe that if the values $\lambda_t(\xi^t)$ are fixed, then (8) decomposes by stage. We thus propose to apply LDR only to the λ variables, leaving the decision variables γ as recourse variables. Imposing the LDR of (9a)

collapses D-MSLP into the following 2SLP, which we refer to as D-LDR-2SLP:

$$\begin{aligned} v^{2S} := \max_{\gamma_1,\alpha} \;& d_1^\top \gamma_1 + \sum_{t\in[T]} \sum_{k\in[K_t]} \mathbb{E}\big[\Phi_{tk}(\xi^t) b_t(\xi^t)^\top\big]\alpha_{tk} + \mathbb{E}[G(\alpha,\xi)] \\ \text{s.t.} \;& \sum_{k\in[K_2]} \mathbb{E}\big[\Phi_{tk}(\xi^2) B_2(\xi^2)^\top\big]\alpha_{2,k} + A_1^\top \alpha_{1,1} + D_1^\top \gamma_1 = h_1, \\ & C_1^\top \alpha_{1,1} + E_1^\top \gamma_1 = c_1, \\ & \gamma_1 \in \mathbb{R}^{n_1}_+, \\ & \alpha_{tk} \in \mathbb{R}^{m_t}, \quad t \in [T],\ k \in [K_t], \end{aligned}$$

where we have dropped the dependence on $\xi_1 = 1$ on the first-stage decision variables. Here, $G(\alpha, \xi)$ is the second-stage value function,

$$G(\alpha,\xi) := \sum_{t\in[2,T]} g_t(\alpha, \xi^t) \tag{10}$$

where for each $t \in [2, T]$,

$$g_t(\alpha,\xi^t) := \max_{\gamma_t} \; d_t(\xi^t)^\top \gamma_t(\xi^t) \tag{11a}$$

$$\text{s.t. } D_t(\xi^t)^\top \gamma_t(\xi^t) = h_t(\xi^t) - A_t(\xi^t)^\top \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\alpha_{tk} - \sum_{k\in[K_{t+1}]} \mathbb{E}\Big[\Phi_{tk}(\xi^{t+1}) B_{t+1}(\xi^{t+1})^\top \Big| \xi^t\Big]\alpha_{t+1,k}, \tag{11b}$$

$$E_t(\xi^t)^\top \gamma_t(\xi^t) = c_t(\xi^t) - C_t(\xi^t)^\top \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\alpha_{tk}, \tag{11c}$$

$$\gamma_t(\xi^t) \in \mathbb{R}^{n_t}_+. \tag{11d}$$

The following proposition, immediate from the definitions of the associated problems, summarizes the relationship between the optimal values of MSLP, D-LDR-LP, and D-LDR-2SLP.

Proposition 3.1. *The following inequalities hold:*

$$z^{MSLP} \geq v^{2S} \geq v^{LDR}.$$

As D-LDR-2SLP is a 2SLP, the discussion in Section 2.3 of methods to obtain an approximate solution to LDR-2SLP applies also to D-LDR-2SLP. In particular, one possibility is to obtain an i.i.d. sample $\{\xi_j\}_{j=1}^N$ of ξ and solve the SAA problem:

$$\hat{v}_N^{2S} := \max_{\gamma_1,\alpha} \; d_1^\top \gamma_1 + \sum_{t\in[T]} \sum_{k\in[K_t]} \mathbb{E}\big[\Phi_{tk}(\xi^t) b_t(\xi^t)^\top\big]\alpha_{tk} + \frac{1}{N}\sum_{j\in[N]} G(\alpha,\xi_j) \tag{12a}$$

$$\text{s.t.} \quad \sum_{k\in[K_2]} \mathbb{E}\big[\Phi_{tk}(\xi^2) B_2(\xi^2)^\top\big]\alpha_{2,k} + A_1^\top \alpha_{1,1} + D_1^\top \gamma_1 = h_1, \tag{12b}$$

$$C_1^\top \alpha_{1,1} + E_1^\top \gamma_1 = c_1, \tag{12c}$$

$$\gamma_1 \in \mathbb{R}^{n_1}_+, \tag{12d}$$

$$\alpha_{tk} \in \mathbb{R}^{m_t}, \quad t \in [T],\ k \in [K_t]. \tag{12e}$$

As in the primal, note that the SAA problem can be solved by decomposition algorithms as the second-stage problem decomposes by both scenario and by stage due to the relationship (10).

3.3 Obtaining lower bounds on z^{MSLP}

Next, we discuss how to use an approximate solution of D-LDR-2SLP to estimate a lower bound on z^{MSLP}. As in the primal case, in order to assure that we obtain a two-stage LDR policy that is feasible for all possible realizations, we consider the possibility of adding a set of constraints $\alpha \in \mathcal{A} \subseteq \mathbb{R}^{\tau_D}$ to D-LDR-2SLP and its SAA counterpart (12) where $\tau_D := \sum_{t\in[T]} K_t m_t$. The following assumption on $\mathcal{A}$ assures that the problem D-LDR-2SLP has relatively complete recourse when the constraints $\alpha \in \mathcal{A}$ are enforced.

Assumption 3.2. *For all $\alpha \in \mathcal{A}$, $G(\alpha,\xi) > -\infty$ $\mathbb{P}$-a.s..*

The following assumption provides a sufficient condition under which the set $\mathcal{A} = \mathbb{R}^{\tau_D}$ satisfies Assumption 3.2 (i.e., no additional constraints are necessary).

Assumption 3.3. *The set $Q_t(\xi^t) := \{x_t \in \mathbb{R}^{p_t},\ s_t \in \mathbb{R}^{q_t} : D_t(\xi^t)s_t + E_t(\xi^t)x_t \geq d_t(\xi^t)\}$ is bounded for all $t \in [T]$ and $\mathbb{P}$ almost all $\xi \in \Xi$.*

A special case of this assumption occurs when x and s variables have explicit upper and lower bounds. An important feature of this assumption is that the sets $Q_t(\xi^t)$ are not required to be uniformly bounded. For example, bounds on the decision variables of the

form $0 \leq x_t(\xi^t) \leq M(\xi^t)$ (and similarly for s variables), are sufficient for satisfying this assumption, even if $M(\xi^t)$ is not bounded over $\xi \in \Xi$.

Proposition 3.4. *If Assumption 3.3 is satisfied, then $\mathcal{A} = \mathbb{R}^{\tau_D}$ satisfies Assumption 3.2.*

Proof. We show that for any given $\alpha \in \mathbb{R}^{\tau_D}$ and $\xi \in \mathbb{R}^{\ell^T}$, (11) is feasible for any $t \in [T]$. Let $R_b(\xi^t)$ and $R_c(\xi^t)$ denote the right-hand sides of the constraints (11b) and (11c), respectively. Then, the dual of (11)

$$\begin{aligned} \min \;\; & R_b(\xi^t)^\top s_t(\xi^t) + R_c(\xi^t)^\top x_t(\xi^t) \\ \text{s.t.} \;\; & D_t(\xi^t)s_t(\xi^t) + E_t(\xi^t)x_t(\xi^t) \geq d_t(\xi^t), \\ & x_t(\xi^t) \in \mathbb{R}^{p_t}, \; s_t(\xi^t) \in \mathbb{R}^{q_t}, \end{aligned}$$

is bounded due to Assumption 3.3. It is also feasible as MSLP is assumed to be feasible. This implies that (11) cannot be infeasible. □

Now, suppose we have an approximate solution $(\hat{\gamma}_1, \hat{\alpha})$ to D-LDR-2SLP, where $\hat{\alpha} \in \mathcal{A}$ for some set $\mathcal{A}$ that satisfies Assumption 3.2. In this case, $(\hat{\gamma}_1, \hat{\alpha})$ defines a feasible solution to D-LDR-2SLP, and hence its objective value provides a lower bound on v^{2S}, and hence is a lower bound on z^{MSLP}. The objective value of $(\hat{\gamma}_1, \hat{\alpha})$ can be estimated by generating an independent sample of ξ, say $\{\xi_j\}_{j=1}^{N'}$ (where possibly $N' \gg N$), and computing

$$d_1^\top \hat{\gamma}_1 + \sum_{t\in[T]} \sum_{k\in[K_t]} \mathbb{E}\big[\Phi_{tk}(\xi^t) b_t(\xi^t)^\top\big] \hat{\alpha}_{tk} + \frac{1}{N'} \sum_{j\in[N']} G(\hat{\alpha}, \xi_j).$$

(The second term in the objective can also be estimated by sampling if the $\mathbb{E}[\cdot]$ terms cannot be computed efficiently.) The values $G(\hat{\alpha}, \xi_j)$ for $j \in [N']$ can also be used to construct a confidence interval on the objective value of $(\hat{\gamma}_1, \hat{\alpha})$, and hence a statistical lower bound on z^{MSLP}.

We close this section by discussing an approach for estimating the gap between a primal two-stage LDR policy defined by $(\hat{x}_1, \hat{\beta})$ and a dual two-stage LDR policy defined by $(\hat{\gamma}_1, \hat{\alpha})$. Following [33], the motivation is that if the same sample (common random numbers) is used in estimating the upper and lower bounds, then the variance of the gap estimator can be reduced if the upper and lower bound sample estimates are positively correlated.

Specifically, given a sample $\{\xi_j\}_{j=1}^{N'}$, the gap observations are then calculated as

$$\begin{aligned}\mathrm{Gap}_j = &\Big[c_1^\top \hat{x}_1 + \sum_{t\in[T]}\sum_{k\in[K_t]} \mathbb{E}\big[\Phi_{tk}(\xi^t)h_t(\xi^t)^\top\big]\hat{\beta}_{tk} + F(\hat{\beta},\xi_j)\Big] \\ &- \Big[d_1^\top \hat{\gamma}_1 + \sum_{t\in[T]}\sum_{k\in[K_t]} \mathbb{E}\big[\Phi_{tk}(\xi^t)b_t(\xi^t)^\top\big]\hat{\alpha}_{tk} + G(\hat{\alpha},\xi_j)\Big],\end{aligned}$$

for $j \in [N']$. These values can then be used to construct a confidence interval on the gap.

4 Illustrative example: capacity expansion

In this section, we present results from a capacity expansion problem which illustrate that two-stage LDR has the potential to yield better bounds and policies for MSLP than static LDR.

4.1 Problem description

We consider a variant of the stochastic capacity expansion problem given in [14]. We wish to determine an investment schedule over T stages for the installation of new capacities of I different power generation technologies, together with some operational decisions to meet demand for power over time. The demand is modeled by a load duration curve, which is approximated by partitioning each stage into J segments (of possibly different length). The demand corresponding to segment $j \in [J]$ in $t \in [T]$ is denoted by d_{tj}. The amount of new capacity of technology $i \in [I]$ added in stage $t \in [T]$ is represented by u_{ti}^+, and is assumed to be available for use immediately, i.e., at the beginning of stage t. The unit cost of u_{ti}^+ is denoted by $c_{ti}^{u^+}$. The state variable s_{ti} represents the current installed capacity of technology $i \in [I]$ in the beginning of stage $t \in [T]$, which incurs holding cost of c_{ti}^s per unit. We assume that it is possible to discard (i.e., remove) some capacity of $i \in [I]$ in $t \in [T]$, denoted by u_{ti}^-, at a (possibly zero) unit cost of $c_{ti}^{u^-}$. The operating level of $i \in [I]$ at $t \in [T]$ for meeting the demand in segment $j \in [J]$ is represented by the decision variable x_{tij}, while the amount of unsatisfied demand is represented as z_{tj}, whose unit costs are c_{tij}^x and c_{tj}^z, respectively. Then, the stochastic capacity expansion problem is formulated as an

MSLP as follows:

$$\min \ \mathbb{E} \sum_{t\in[T]} \left[\sum_{i\in[I]} \left(c^{u^+}_{ti} u^+_{ti}(\xi^t) + c^{u^-}_{ti} u^-_{ti}(\xi^t) + c^s_{ti}\, s_{ti}(\xi^t) + \sum_{j\in[J]} c^x_{tij} x_{tij}(\xi^t) \right) + \sum_{j\in[J]} c^z_{tj} z_{tj}(\xi^t) \right] \tag{13a}$$

$$\text{s.t.} \ \ s_{ti}(\xi^t) - s_{t-1,i}(\xi^{t-1}) - u^+_{ti}(\xi^t) \ + u^-_{ti}(\xi^t) = 0, \ t \in [T], \ \mathbb{P}\text{-a.s.}, \ i \in [I], \tag{13b}$$

$$s_{ti}(\xi^t) - x_{tij}(\xi^t) \geq 0, \qquad t \in [T], \ \mathbb{P}\text{-a.s.}, \ i \in [I], \ j \in [J], \tag{13c}$$

$$\sum_{i\in[I]} x_{tij}(\xi^t) + z_{tj}(\xi^t) \geq d_{tj}(\xi^t), \qquad t \in [T], \ \mathbb{P}\text{-a.s.}, \ j \in [J], \tag{13d}$$

$$0 \leq z_{tj}(\xi^t) \leq d_{tj}(\xi^t), \qquad t \in [T], \ \mathbb{P}\text{-a.s.}, \ j \in [J], \tag{13e}$$

$$0 \leq u^+_{ti}(\xi^t) \leq M^{u^+}_{ti}(\xi^t), \qquad t \in [T], \ \mathbb{P}\text{-a.s.}, \ i \in [I], \tag{13f}$$

$$0 \leq u^-_{ti}(\xi^t) \leq M^{u^-}_{ti}(\xi^t), \qquad t \in [T], \ \mathbb{P}\text{-a.s.}, \ i \in [I], \tag{13g}$$

$$0 \leq s_{ti}(\xi^t) \leq M^s_{ti}(\xi^t), \qquad t \in [T], \ \mathbb{P}\text{-a.s.}, \ i \in [I], \tag{13h}$$

$$x_{tij}(\xi^t) \geq 0, \qquad t \in [T], \ \mathbb{P}\text{-a.s.}, \ i \in [I], \ j \in [J], \tag{13i}$$

The objective function (13a) minimizes the expected total cost. The constraints (13b) are the only state equations, which keep track of the available capacity of each technology. Constraints (13c) limit the operating levels to the available capacity level, while (13d) ensure that either demand is met, or unmet demand is recorded in the z_{tj} variable values. Constraints (13e)-(13h) represent the bounds on the shortfall, installation, removal, and inventory level variables, respectively. Note that (13h) constitute upper bounds also on the x variables due to (13c), and thus this formulation satisfies Assumption 3.3. In addition, we assume that $M^{u^-}_{ti}(\xi^t) \geq M^s_{t-1,i}(\xi^{t-1})$ which ensures that this formulation satisfies relatively complete recourse, Assumption 2.3, since at stage t, given any feasible value of $s_{t-1,i}(\xi^{t-1})$, it is feasible to set $u^-_{ti}(\xi^t) = s_{t-1,i}(\xi^{t-1})$ for $i \in [I]$, $z_{tj}(\xi^t) = d_{tj}(\xi^t)$ for $j \in [J]$ and all remaining variables to zero.

The dual of (13) is provided in (15) in Appendix B.

4.2 Data

To the extent possible, we use data from [14], which focuses on a German system, although we extend their 3-stage example to $T = 3, \ldots, 10$. There are $I = 3$ technologies (coal-fired power plant, combined cycle gas turbine and open cycle gas turbine). In [14] each stage is divided into $L = 8$ periods, and $W = 5$ wind regimes are considered for each period. We

model this as $J = LW = 40$ segments at each stage, corresponding to each period/wind regime pair. For $t \geq 2$, the demand corresponding to the segment $j = (\ell, w) \in [L] \times [W]$ is modeled as

$$d_{tj}(\xi^t) = \max\left\{ d_{0,\ell} \prod_{r=2}^{t} \xi_r^g - \eta_w K_t^w \prod_{r=2}^{t} \xi_r^w, 0 \right\}$$

where $d_{0,\ell}$ is the base demand value of period ℓ, ξ_t^g is a random variable reflecting the demand growth of stage t, η_w is the parameter denoting the wind efficiency, K_t^w is the wind power generation target, and ξ_t^w is a random variable representing the growth in the wind power generation in stage t. The values of $d_{0,\ell}$ and η_w are from Tables 2 and 3 of [14], and are reproduced in Appendix A. We use $K_2^w = 36.64$ and $K_t^w = 45.75$ for all $t \geq 3$. We assume ξ_t^g has lognormal distribution with $\mu = 0.2$ and $\sigma = 0.1 + 0.01t$, and ξ_t^w has lognormal distribution with $\mu = 0.15$ and $\sigma = 0.25 + 0.025t$. For the first stage, we use the deterministic demand values of $d_{1,j=(\ell,w)} = d_{0,\ell}\mathbb{E}[\xi_1^g] - \eta_w \mathbb{E}[\xi_1^w] = 1.229 d_{0,\ell} - 1.207\eta_w$. The units of all demands (and all primal decision variables) are gigawatts.

We assume there are no holding costs and no costs for removing capacity, i.e., we use $c_{ti}^{u^-} = c_{ti}^s = 0$ for all $i \in [I], t \in [T]$. We use discounting to determine the other costs, setting $c_{ti}^{u^+} = 5\iota_i/1.1^t$, $c_{ti,j=(\ell,w)}^x = 0.001 c_i \tau_\ell \tau_w / 1.1^t$ and $c_{t,j=(\ell,w)}^z = \tau_\ell \tau_w / 1.1^t$ where the values of the annualized costs ι_i, operation costs c_i, τ_ℓ and τ_w values are from [14] (see Appendix A). All costs are in million of Euros.

Lastly, for the upper bounds in the primal model, we use

$$M_{ti}^s(\xi^t) = M_{ti}^{u^+}(\xi^t) = M \sum_{k \in [K_t]} \Phi_{tk}(\xi^t) \text{ and } M_{ti}^{u^-}(\xi^t) = M_{t-1,i}^s(\xi^{t-1}),$$

where in our experiments $M = 100$.

4.3 Implementation details

We compare the primal and dual bounds obtained using two-stage and static LDR. We choose the LDR basis functions to be the underlying random variables that determine the random demands. Specifically, for any $t \in [T]$, we let $K_t = 1 + 2(t-1)$, $\Phi_{t1}(\xi^t) = 1$, and

$$\Phi_{tk}(\xi^t) = \begin{cases} \prod_{r=2}^{k} \xi_r^g, & k = 2, \ldots, t \\ \prod_{r=2}^{k-t+1} \xi_r^w, & k = t+1, \ldots, 2t-1. \end{cases}$$

Because assumptions A3 and A4 do not hold for this problem (the random variables do not have bounded support and $\mathbb{E}(\xi|\xi^t)$ is not linear in ξ^t), the reformulation approach from [12, 29, 50] used for static LDR cannot be applied to solve LDR-LP and D-LDR-LP. We therefore use a sampling strategy to approximately solve these problems. Specifically, the sample approximations of LDR-LP and D-LDR-LP are identical to LDR-LP and D-LDR-LP, respectively, except that the infinite set of constraints $\mathbb{P}$-a.s. are replaced by the finite set corresponding to the sample. Approximate solutions to LDR-2SLP and D-LDR-2SLP are obtained by solving the SAA problems (5) and (12), respectively. We solve all sample approximations using the same sample of size $N = 100T$. We then use an independent sample of size $N' = 10000T$ to obtain an unbiased estimate of the solution quality, and again use the same sample in all evaluations.

Models LDR-2SLP and D-LDR-2SLP corresponding to model (13) and its dual are given in Appendix B. Although the MSLP given in (13) has relatively complete recourse (Assumption 2.3), its two-stage primal approximation LDR-2SLP does not have relatively complete recourse. The reason is that the LDR variables corresponding to s_{ti} for $i \in [I],\ t \in [2,T]$, namely $\{\beta_{tki}\}_{k\in[K_t]}$, have no bounds, and therefore there exist values of these variables that would make the recourse problem infeasible under some scenarios. However, as shown in Proposition B.1 in Appendix B, if we define $\mathcal{B} := [0,M]^{\tau_P}$, then $\mathcal{B}$ satisfies Assumption 2.2. We therefore add the constraints $\beta \in \mathcal{B}$ to the sample-average approximation of LDR-2SLP. In order to improve the chances that the policy obtained by sample approximation of the static LDR model, LDR-LP, is feasible, we also introduce these constraints in the sample approximation of LDR-LP.

With the addition of the constraints $\beta \in \mathcal{B}$, any solution to the SAA problem (5) provides a feasible solution to LDR-2SLP, and hence defines a feasible policy to MSLP and using the N' independent replications it can be used to derive a statistical upper bound on z^{MSLP}. Similarly, as Assumption 3.3 is satisfied, any solution to the SAA problem (12) provides a feasible solution to D-LDR-2SLP, and hence evaluating this solution using N' independent replications yields a statistical lower bound on z^{MSLP}. Unfortunately, the corresponding sample approximations of LDR-LP and D-LDR-LP do not yield policies (primal or dual) that are feasible under all scenarios. Thus, when evaluating these policies with the independent replications, we report two measures: the average objective value over scenarios that are feasible, and the fraction of scenarios that are infeasible. By averaging only over feasible scenarios, these estimates are optimistically biased, i.e., they underestimate the bound on the primal problem, and overestimate the bound for the dual

problem. As a result, these estimates do not necessarily provide valid (statistical) upper and lower bounds on z^{MSLP}, but we use them to provide a "best case" estimate when comparing to the estimates obtained from the two-stage LDR policies.

All of our numerical results are carried out using IBM ILOG CPLEX 12.6 as the LP solver on a Linux workstation with 24, 2.67 GHz Intel Xeon CPUs and 132 GB memory. Default CPLEX settings are used, except the optimality tolerance (EpOpt parameter) is set to 10^{-9} for better accuracy. As the focus of this numerical example is to compare the quality of solutions, we made no effort to make our implementation efficient. In particular, the SAA problems are solved in the extensive form (i.e., not via decomposition or a cutting plane algorithm), and when evaluating solutions each optimization model is re-built from scratch, requiring significant time in model building and deleting that could be avoided with an improved implementation. The time to solve the sample approximations of LDR-2SLP and D-LDR-2SLP ranged between two minutes and five hours. The solution times of the sample approximation of LDR-LP was comparable to that of LDR-2SLP, while solving the sample approximation of D-LDR-LP took from ten minutes up to a day. The time spent evaluating the solutions on the large sample of $N' = 10000T$ scenarios ranged between five hours and a day, for all methods. However, for each instance, the total time spent by CPLEX actually solving the optimization models over all these scenarios was less than 30 minutes, indicating that the evaluation time could be very significantly reduced with a better implementation.

4.4 Upper bound comparison

We start by comparing the expected costs of two-stage and static LDR policies. Table 1 presents 95% confidence intervals (CIs) on expected cost of the two-stage and static LDR policies, represented by their means and widths (CI $\pm$). For the two-stage LDR policy, the upper end of the CI is a statistical upper bound on z^{MSLP}. We also report under column 'Inf. (%)' the percentage of the scenarios (out of $10000T$ evaluated scenarios) for which the static LDR policy is infeasible. To give an idea of the relative improvement in the expected policy cost obtained by using the two-stage LDR, the column 'CI Ratio', presents the ratio of the CI mean of the two-stage LDR policy to the CI mean of the static LDR policy. We observe that the expected cost of the two-stage LDR policy is between 20% and 34% lower than that of the static LDR policy, with the most significant differences occurring with larger time stages. We also observe that the static LDR policy is frequently infeasible

	Two-stage LDR		Static LDR			
T	CI Mean	CI ±	CI Mean	CI ±	Inf. (%)	CI Ratio
3	119874	173	149165	113	2.1	0.80
4	150785	267	200339	192	2.3	0.75
5	186915	374	256019	280	3.5	0.73
6	228791	495	323645	394	3.5	0.71
7	276320	630	397723	529	5.0	0.69
8	330223	794	482306	685	5.1	0.68
9	392097	986	588358	909	3.7	0.67
10	463083	1216	701287	1156	3.5	0.66

Table 1: Confidence intervals for expected costs of the primal policies

when evaluated under the $10000T$ independent scenarios.

4.5 Lower bound comparison and gap estimates

Next, we compare the lower bounds obtained from using two-stage and static LDR policies in the dual of (13). Table 2 presents the CI means and widths for both approaches, the percentage of infeasible scenarios (out of $10000T$) for the static LDR policy, and the ratio of the CI mean of the two-stage LDR policy to that of the static LDR policy. We find that

	Two-stage LDR		Static LDR			
T	CI Mean	CI ±	CI Mean	CI ±	Inf. (%)	CI Ratio
3	117617	233	117475	184	8.2	1.001
4	148156	300	147245	257	8.6	1.006
5	183405	394	180501	348	10.2	1.016
6	223649	637	217541	458	15.0	1.028
7	267633	1119	259015	571	18.2	1.033
8	313323	2090	305156	723	22.1	1.027
9	366025	3330	357417	888	23.1	1.024
10	426437	5342	421358	1111	24.9	1.012

Table 2: Confidence intervals for expected costs of the dual policies

the obtained bounds are similar, although the bound from the two-stage LDR policy is between 0.1% to 3.3% better than that from the static LDR policy. We also find that the approximate static LDR policy is very often infeasible when evaluated on an independent

sample, and hence the CI Mean estimate for that policy (which excludes the infeasible scenarios) may be significantly higher than the true optimal value of of D-LDR-LP.

Finally, we present in Table 3 estimates of the gap between the expected cost of the primal and dual two-stage and static LDR policies. These gap estimates are obtained using the procedure discussed at the end of Section 3.3. For the static LDR, scenarios for which the static LDR policy was infeasible in either the primal or dual problem are excluded from the CI calculation, and the column 'Inf. (%)' reports the percentage of excluded scenarios. In addition to reporting the mean (CI Mean) and width (CI $\pm$) of the confidence intervals, we also report under 'Gap (%)' the ratio of the gap CI width to the mean of the upper bound estimate, which gives an estimate of the relative optimality gap between the primal and dual policies. We see that that the two-stage LDRs yield significant reduction in gap over the static LDRs. The gaps are smaller for instances with fewer stages T, but never exceed 10%.

	Two-stage LDR			Static LDR			
T	CI Mean	CI $\pm$	Gap (%)	CI Mean	CI $\pm$	Gap (%)	Inf. (%)
3	2257	81	2.0	32027	106	21.5	8.3
4	2629	48	1.8	53760	140	26.9	9.3
5	3511	50	1.9	76338	187	29.9	11.0
6	5142	383	2.4	105906	243	32.8	15.4
7	8687	905	3.5	137943	302	34.8	19.0
8	16900	1908	5.7	174842	365	36.3	22.7
9	26072	3155	7.5	223454	465	38.1	23.5
10	36646	5175	9.0	271114	573	38.7	25.3

Table 3: Optimality gap estimates

5 Concluding remarks

We propose two-stage LDRs, a new approximate solution method for MSLPs. This approach has two advantages. Due to the flexibility in the recourse decisions, our method potentially yields better (at least not worse) bounds and policies than standard static LDR policies. In addition, as our approach is based on sampling and 2SLP, it works with very mild assumptions and can take advantage of existing literature on methods for approximately solving 2SLP problems. We illustrate the new approach on a capacity expansion

example, which indicates that two-stage LDR policies have potential to yield significantly better policies than static LDR policies.

In future research it will be interesting to test the use of two-stage LDR on more problems, and to investigate if there are problem classes where two-stage LDR yields policies that are provably optimal or near-optimal.

As we mentioned in the introduction, in the primal problem, two-stage LDR can be directly applied to multi-stage stochastic *mixed integer programs*, provided the integrality restrictions are imposed only on the recourse variables. Since availability of algorithms for multi-stage stochastic mixed integer programs is very limited, it will be interesting to explore this extension further, in particular possibly using ideas from [6] to obtain a decision rule in the case the state variables also have integrality constraints.

Acknowledgements. This work is supported in part by the National Science Foundation under grant CMMI-1634597, and by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under contract number DE-AC02-06CH11357.

References

[1] S. Ahmed. Multistage stochastic optimization. `https://www.ima.umn.edu/materials/2015-2016/ND8.1-12.16/2` New Directions Short Course on Mathematical Optimization, 2016.

[2] I. Bakir, N. Boland, B. Dandurand, and A. Erera. Scenario set partition dual bounds for multistage stochastic programming: A hierarchy of bounds and a partition sampling approach, 2016. `http://www.optimization-online.org/DB_FILE/2016/01/5311.pdf`.

[3] D. Bampou and D. Kuhn. Scenario-free stochastic programming with polynomial decision rules. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pages 7806–7812. IEEE, 2011.

[4] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Math. Program.*, 99(2):351–376, 2004.

[5] D. Bertsimas and C. Caramanis. Adaptability via sampling. In *Decision and Control, 2007 46th IEEE Conference on*, pages 4717–4722. IEEE, 2007.

[6] D. Bertsimas and A. Georghiou. Design of near optimal decision rules in multistage adaptive mixed-integer optimization. *Oper. Res.*, 63(3):610–627, 2015.

[7] J.R. Birge. Aggregation bounds in stochastic linear programming. *Math. Program.*, 31(1):25–41, 1985.

[8] J.R. Birge. Decomposition and partitioning methods for multistage stochastic linear programs. *Oper. Res.*, 33(5):989–1007, 1985.

[9] J.R. Birge, C.J. Donohue, D.F. Holmes, and O.G. Svintsitski. A parallel implementation of the nested decomposition algorithm for multistage stochastic linear programs. *Math. Program.*, 75(2):327–352, 1996.

[10] G.C. Calafiore and M.C. Campi. The scenario approach to robust control design. *IEEE Trans. Automat. Control*, 51:742–753, 2006.

[11] M.S. Casey and S. Sen. The scenario generation algorithm for multistage stochastic linear programming. *Math. Oper. Res.*, 30(3):615–631, 2005.

[12] X. Chen, M. Sim, P. Sun, and J. Zhang. A linear decision-based approximation approach to stochastic programming. *Oper. Res.*, 56(2):344–357, 2008.

[13] Z.-L. Chen and W.B. Powell. Convergent cutting-plane and partial-sampling algorithm for multistage stochastic linear programs with recourse. *J. Optim. Theory Appl.*, 102(3):497–524, 1999.

[14] G. de Maere d'Aertrycke, A. Shapiro, and Y. Smeers. Risk exposure and lagrange multipliers of nonanticipativity constraints in multistage stochastic problems. *Math. Meth. Oper. Res.*, 77(3):393–405, 2013.

[15] M. Dyer and L. Stougie. Computational complexity of stochastic programming problems. *Math. Program.*, 106(3):423–432, 2006.

[16] M.J. Eisner and P. Olsen. Duality for stochastic programming interpreted as L.P. in L_p-space. *SIAM J. Appl. Math.*, 28(4):779–792, 1975.

[17] M.C. Fu. Feature article: Optimization for simulation: Theory vs. practice. *INFORMS Journal on Computing*, 14(3):192–215, 2002.

[18] S.J. Garstka and R.J.B. Wets. On decision rules in stochastic programming. *Math. Program.*, 7(1):117–143, 1974.

[19] H.I. Gassmann. Mslip: A computer code for the multistage stochastic linear programming problem. *Math. Program.*, 47(1):407–423, 1990.

[20] A. Georghiou, W. Wiesemann, and D. Kuhn. Generalized decision rule approximations for stochastic programming via liftings. *Math. Program.*, 152(1-2):301–338, 2015.

[21] G.A. Hanasusanto, D. Kuhn, and W. Wiesemann. A comment on "computational complexity of stochastic programming problems". *Math. Program.*, 159(1):557–569, 2016.

[22] H. Heitsch and W. Römisch. Scenario tree modeling for multistage stochastic programs. *Math. Program.*, 118(2):371–406, 2009.

[23] J.L. Higle and S. Sen. Multistage stochastic convex programs: Duality and its implications. *Ann. Oper. Res.*, 142(1):129–146, 2006.

[24] T. Homem de Mello. On rates of convergence for stochastic optimization problems under non-independent and identically distributed sampling. *SIAM J. Optim.*, 19:524–551, 2008.

[25] L.J. Hong and B.L. Nelson. A brief introduction to optimization via simulation. In *Winter Simulation Conference*, WSC '09, pages 75–85. Winter Simulation Conference, 2009.

[26] K. Høyland and S.W. Wallace. Generating scenario trees for multistage decision problems. *Management Sci.*, 47(2):295–307, 2001.

[27] G. Infanger and D.P. Morton. Cut sharing for multistage stochastic linear programs with interstage dependency. *Math. Program.*, 75(2):241–256, 1996.

[28] M. Koivu. Variance reduction in sample approximations of stochastic programs. *Math. Program.*, 103:463–485, 2005.

[29] D. Kuhn, W. Wiesemann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Math. Program.*, 130(1):177–209, 2011.

[30] C. Lemaréchal, A. Nemirovskii, and Y. Nesterov. New variants of bundle methods. *Math. Program.*, 69:111–147, 1995.

[31] J. Linderoth, A. Shapiro, and S. Wright. The empirical behavior of sampling methods for stochastic programming. *Ann. Oper. Res.*, 142:215–241, 2006.

[32] K. Linowsky and A.B. Philpott. On the convergence of sampling-based decomposition algorithms for multistage stochastic programs. *J. Optim. Theory Appl.*, 125(2):349–366, 2005.

[33] W.-K. Mak, D.P. Morton, and R.K. Wood. Monte Carlo bounding techniques for determining solution quality in stochastic programs. *Oper. Res. Lett.*, 24:47–56, 1999.

[34] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.

[35] T. Pennanen. Epi-convergent discretizations of multistage stochastic programs via integration quadratures. *Math. Program.*, 116:461–479, 2009.

[36] M.V.F. Pereira and L.M.V.G. Pinto. Multi-stage stochastic optimization applied to energy planning. *Math. Program.*, 52(1-3):359–375, 1991.

[37] A.B. Philpott and V.L. De Matos. Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion. *European J. Oper. Res.*, 218(2):470–483, 2012.

[38] A.B. Philpott and Z. Guan. On the convergence of stochastic dual dynamic programming and related methods. *Oper. Res. Lett.*, 36(4):450–455, 2008.

[39] B.T. Polyak and A.B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM J. Control Optim.*, 30:838–855, 1992.

[40] W.B. Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons, 2007.

[41] H. Robbins and S. Monro. A stochastic approximation method. *Ann. Math. Stat.*, 22:400–407, 1951.

[42] R.T. Rockafellar and R.J.-B. Wets. Scenarios and policy aggregation in optimization under uncertainty. *Math. Oper. Res.*, 16(1):119–147, 1991.

[43] A. Ruszczyński. A regularized decomposition method for minimizing a sum of polyhedral functions. *Mathematical Programming*, 35(3):309–333, 1986.

[44] S. Sen and Z. Zhou. Multistage stochastic decomposition: a bridge between stochastic programming and approximate dynamic programming. *SIAM J. Optim.*, 24(1):127–153, 2014.

[45] A. Shapiro. Stochastic programming approach to optimization under uncertainty. *Math. Program.*, 112:183–220, 2008.

[46] A. Shapiro. Analysis of stochastic dual dynamic programming method. *European J. Oper. Res.*, 209(1):63–72, 2011.

[47] A. Shapiro. Topics in stochastic programming. *CORE Lecture Series, Universite Catholique de Louvain*, 2011.

[48] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on stochastic programming: modeling and theory*, volume 16. SIAM, 2014.

[49] A. Shapiro and T. Homem-de Mello. On the rate of convergence of optimal solutions of monte carlo approximations of stochastic programs. *SIAM J. Optim.*, 11(1):70–86, 2000.

[50] A. Shapiro and A. Nemirovski. On complexity of stochastic programming problems. In *Continuous optimization*, pages 111–146. Springer, 2005.

[51] R.M. Van S. and R. Wets. L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM J. Appl. Math.*, 17(4):638–663, 1969.

[52] S.W. Wallace and W.T. Ziemba. *Applications of stochastic programming*, volume 5. SIAM, 2005.

A Data for the capacity expansion problem

	$i=1$	$i=2$	$i=3$
ι_i (k€/MW)	245.8	113.9	57.8
c_i (€/MWh)	41.9	58.9	90.8

Table 4: Fixed annual cost and operation cost (Table 1 in [14])

	$\ell=1$	$\ell=2$	$\ell=3$	$\ell=4$	$\ell=5$	$\ell=6$	$\ell=7$	$\ell=8$
d_0, ℓ (GW)	77.1	71.4	65.7	60.1	54.4	48.8	43.1	37.4
τ_ℓ (h)	68	677	1585	1781	1367	1688	1289	305

Table 5: Initial demand (Table 2 in [14])

	$w=1$	$w=2$	$w=3$	$w=4$	$w=5$
η_w (%)	92.9	81.8	54.9	21.2	0.0
τ_w (%)	19.8	21.78	18.2	26.7	13.5

Table 6: Wind regimes (Table 3 in [14])

B Additional models for the capacity expansion example

LDR-2SLP of the capacity expansion model is obtained by substituting

$$s_{ti}(\xi^t) = \sum_{k \in [K_i]} \Phi_{tk}(\xi^t)\beta_{tki}$$

in (13). Dropping ξ^t dependences for variables to simplify the notation, we obtain

$$
\begin{aligned}
\min \quad & \sum_{i\in[I]} \Big(c^{u^+}_{1i} u^+_{1i} + c^{u^-}_{1i} u^-_{1i} + \sum_{j\in[J]} c^x_{1ij} x_{1ij}(\xi^t) \Big) + \sum_{j\in[J]} c^z_{1j} z_{1j} \\
& \quad + \sum_{t\in[T]} \sum_{i\in[I]} c^s_{ti} \sum_{k\in[K_t]} \mathbb{E}\big[\Phi_{tk}(\xi^t)\big] \beta_{tki} + \sum_{t\in[2,T]} \mathbb{E}[f_t(\beta, \xi^t)] \\
\text{s.t.} \quad & \beta_{1i1} - u^+_{1i} + u^-_{1i} = 0, && i \in [I], \\
& \beta_{1i1} - x_{1ij} \geq 0, && i \in [I],\ j \in [J], \\
& \sum_{i\in[I]} x_{1ij} + z_{1j} \geq d_{1j}, && j \in [J], \\
& 0 \leq z_{1j} \leq d_{1j}, && j \in [J], \\
& 0 \leq u^+_{1i} \leq M, \quad 0 \leq u^-_{1i} \leq M, \quad 0 \leq \beta_{1i1} \leq M, && i \in [I], \\
& x_{1ij} \geq 0, && i \in [I],\ j \in [J],
\end{aligned}
$$

where, for $t \in [2, T]$, $f_t(\beta, \xi^t)$ is defined as the optimal objective value of the following problem:

$$
\begin{aligned}
\min \quad & \sum_{i\in[I]} \Big(c^{u^+}_{ti} u^+_{ti} + c^{u^-}_{ti} u^-_{ti} + \sum_{j\in[J]} c^x_{tij} x_{tij} \Big) + \sum_{j\in[J]} c^z_{tj} z_{tj} && && \text{(14a)} \\
\text{s.t.} \quad & u^+_{ti} - u^-_{ti} = \sum_{k\in[K_t]} \Phi_{tk}(\xi^t) \beta_{tki} - \sum_{k\in[K_{t-1}]} \Phi_{tk}(\xi^{t-1}) \beta_{t-1,k,i}, && i \in [I], && \text{(14b)} \\
& x_{tij} \leq \sum_{k\in[K_t]} \Phi_{tk}(\xi^t) \beta_{tki}, && i \in [I],\ j \in [J], && \text{(14c)} \\
& \sum_{i\in[I]} x_{tij} + z_{tj} \geq d_{tj}(\xi^t), && j \in [J], && \text{(14d)} \\
& 0 \leq z_{tj} \leq d_{tj}(\xi^t), && j \in [J], && \text{(14e)} \\
& 0 \leq u^+_{ti} \leq M \sum_{k\in[K_t]} \Phi_{tk}(\xi^t), && i \in [I], && \text{(14f)} \\
& 0 \leq u^-_{ti} \leq M \sum_{k\in[K_{t-1}]} \Phi_{tk}(\xi^{t-1}), && i \in [I], && \text{(14g)} \\
& 0 \leq M \sum_{k\in[K_t]} \Phi_{tk}(\xi^t) - \sum_{k\in[K_t]} \Phi_{tk}(\xi^t) \beta_{tki}, && i \in [I], && \text{(14h)} \\
& 0 \leq \sum_{k\in[K_t]} \Phi_{tk}(\xi^t) \beta_{tki}, && i \in [I], && \text{(14i)} \\
& x_{tij} \geq 0, && i \in [I],\ j \in [J], && \text{(14j)}
\end{aligned}
$$

We note that LDR-2SLP of the capacity expansion model does not have relatively complete recourse since the recourse constraints (14h) and (14i) might be violated under some scenarios. However, the following proposition shows that the addition of the constraints $\beta \in [0, M]^{\tau_P}$ makes the problem (14) feasible, thus yields relatively complete recourse.

Proposition B.1. *Let $t \in [T]$. If $0 \le \beta_{tki} \le M$ for all $i \in [I], k \in [K_t]$, then $f_t(\beta, \xi^t) < \infty$ $\mathbb{P}$-a.s..*

Proof. Note because the random data ξ_t^g and ξ_t^w have lognormal distribution, the defined basis functions $\Phi_{tk}(\xi^t)$ are nonnegative $\mathbb{P}$-a.s.. Define a solution $(\hat{u}^+, \hat{u}^-, \hat{x}, \hat{z})$ by setting $\hat{u}_{ti}^+ = \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\beta_{tki}$ and $\hat{u}_{ti}^- = \sum_{k\in[K_{t-1}]} \Phi_{tk}(\xi^{t-1})\beta_{t-1,k,i}$ for all $i \in [I]$, $\hat{x} = 0$ and $\hat{z}_{tj} = d_{tj}(\xi^t)$ for all $j \in [J]$. At $(\hat{u}^+, \hat{u}^-, \hat{x}, \hat{z})$, the constraints (14b), (14d), and (14e) trivially hold. By definition, we have $\Phi_{tk}(\xi^t) \ge 0$ for all $k \in [K_t]$, and together with $\beta_{tki} \ge 0$ for all $i \in [I], k \in [K_t]$, this implies that the constraints (14c), (14f), (14g), and (14i) are satisfied. Lastly, the nonnegativity of the basis function values and the upper bound of M on the β variables imply that (14h) is satisfied. Therefore, $(\hat{u}^+, \hat{u}^-, \hat{x}, \hat{z})$ is feasible to (14) and hence $f_t(\beta, \xi^t) < \infty$. □ □

Let $\lambda, \gamma, \theta^+, \theta^-, \pi^{u^+}, \pi^{u^-}, \pi^s$ be the dual variables associated with the constraints (13b)-(13h) in (13), respectively. Then, the dual of (13) is:

$$
\begin{aligned}
\max\ & \mathbb{E} \sum_{t\in[T]} \Big[\sum_{j\in[J]} d_{tj}(\xi^t)\big(\theta_{tj}^+(\xi^t) + \theta_{tj}^- \ \ (\xi^t)\big) \\
& \quad + \sum_{i\in[I]} \big(M_{ti}^{u^+}(\xi^t)\pi_{ti}^{u^+}(\xi^t) + M_{ti}^{u^-}(\xi^t)\pi_{ti}^{u^-}(\xi^t) + M_{ti}^{s}(\xi^t)\pi_{ti}^{s}(\xi^t)\big)\Big] && \text{(15a)} \\
\text{s.t.}\ & \lambda_{ti}(\xi^t) - \mathbb{E}[\lambda_{t+1,i}(\xi^{t+1}) \mid \xi^t] \\
& \quad + \sum_{j\in[J]} \gamma_{tij}(\xi^t) + \pi_{ti}^s(\xi^t) \le c_{ti}^s, \quad t \in [T],\ \mathbb{P}\text{-a.s.},\ i \in [I], && \text{(15b)} \\
& \pi_{ti}^{u^+}(\xi^t) - \lambda_{ti}(\xi^t) \le c_{ti}^{u^+}, \qquad t \in [T],\ \mathbb{P}\text{-a.s.},\ i \in [I], && \text{(15c)} \\
& \pi_{ti}^{u^-}(\xi^t) + \lambda_{ti}(\xi^t) \le c_{ti}^{u^-}, \qquad t \in [T],\ \mathbb{P}\text{-a.s.},\ i \in [I], && \text{(15d)} \\
& \theta_{tj}^+(\xi^t) + \theta_{tj}^-(\xi^t) \ \le c_{tj}^z, \qquad t \in [T],\ \mathbb{P}\text{-a.s.},\ j \in [J], && \text{(15e)} \\
& \theta_{tj}^+(\xi^t) - \gamma_{tij}(\xi^t) \le c_{tij}^x, \qquad t \in [T],\ \mathbb{P}\text{-a.s.},\ i \in [I],\ j \in [J], && \text{(15f)} \\
& \gamma_{tij}(\xi^t) \ge 0, \qquad t \in [T],\ \mathbb{P}\text{-a.s.},\ i \in [I],\ j \in [J], && \text{(15g)} \\
& \theta_{tj}^+(\xi^t) \ge 0, \ \ \theta_{tj}^-(\xi^t) \le 0 \qquad t \in [T],\ \mathbb{P}\text{-a.s.},\ j \in [J], && \text{(15h)} \\
& \pi_{ti}^{u^+}(\xi^t), \pi_{ti}^{u^-}(\xi^t), \pi_{ti}^s(\xi^t) \le 0 \qquad t \in [T],\ \mathbb{P}\text{-a.s.},\ i \in [I]. && \text{(15i)}
\end{aligned}
$$

D-LDR-2SLP of the capacity expansion model is obtained by substituting

$$\lambda_{ti}(\xi^t) = \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\alpha_{tki}$$

in (15). Dropping ξ^t dependences for variables to simplify the notation, we obtain

$$
\begin{aligned}
\max \quad & \sum_{j\in[J]} d_{1j}\big(\theta^+_{1j}+\theta^-_{1j}\big) + \sum_{i\in[I]} M\big(\pi^{u^+}_{1i}+\pi^{u^-}_{1i}+\pi^s_{1i}\big) + \sum_{t\in[2,T]} \mathbb{E}[g_t(\alpha,\xi^t)] \\
\text{s.t.} \quad & \alpha_{1i1} - \sum_{k\in[K_2]} \mathbb{E}\big[\Phi_{tk}(\xi^2)\big]\alpha_{2,i,k} + \sum_{j\in[J]} \gamma_{1ij} + \pi^s_{1i} \le c^s_{1i}, && i\in[I], \\
& \pi^{u^+}_{1i} - \alpha_{1i1} \le c^{u^+}_{1i}, && i\in[I], \\
& \pi^{u^-}_{1i} + \alpha_{1i1} \le c^{u^-}_{1i}, && i\in[I], \\
& \theta^+_{1j} + \theta^-_{1j} \le c^z_{1j}, && j\in[J], \\
& \theta^+_{1j} - \gamma_{1ij} \le c^x_{1ij}, && i\in[I],\ j\in[J] \\
& \gamma_{1ij} \ge 0, && i\in[I],\ j\in[J] \\
& \theta^+_{1j} \ge 0, \quad \theta^-_{1j} \le 0 && j\in[J], \\
& \pi^{u^+}_{1i}, \pi^{u^-}_{1i}, \pi^s_{1i} \le 0 && i\in[I],
\end{aligned}
$$

where, for $t\in[2,T]$, $g_t(\alpha,\xi^t)$ is defined as the optimal objective value of the following problem:

$$
\begin{aligned}
\max \quad & \sum_{j\in[J]} d_{tj}(\xi^t)\big(\theta^+_{tj}+\theta^-_{tj}\big) + \sum_{i\in[I]} \Big(M \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\Big)\big(\pi^{u^+}_{ti}+\pi^s_{ti}\big) \\
& \qquad\qquad + \Big(M \sum_{k\in[K_{t-1}]} \Phi_{tk}(\xi^{t-1})\Big)\pi^{u^-}_{ti} \\
\text{s.t.} \quad & \sum_{j\in[J]} \gamma_{tij} + \pi^s_{ti} \le c^s_{ti} - \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\alpha_{tki} \\
& \qquad\qquad + \sum_{k\in[K_{t+1}]} \mathbb{E}\Big[\Phi_{tk}(\xi^{t+1})\Big|\xi^t\Big]\alpha_{t+1,i,k}, && i\in[I], \\
& \pi^{u^+}_{ti} \le c^{u^+}_{ti} + \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\alpha_{tki}, && i\in[I], \\
& \pi^{u^-}_{ti} \le c^{u^-}_{ti} - \sum_{k\in[K_t]} \Phi_{tk}(\xi^t)\alpha_{tki}, && i\in[I], \\
& \theta^+_{tj} + \theta^-_{tj} \le c^z_{tj}, && j\in[J], \\
& \theta^+_{tj} - \gamma_{tij} \le c^x_{tij}, && i\in[I],\ j\in[J], \\
& \gamma_{tij} \ge 0, && i\in[I],\ j\in[J], \\
& \theta^+_{tj} \ge 0, \quad \theta^-_{tj} \le 0 && j\in[J], \\
& \pi^{u^+}_{ti}, \pi^{u^-}_{ti}, \pi^s_{ti} \le 0 && i\in[I].
\end{aligned}
$$